arXiv:2006.08482v1 [stat.ME] 15 Jun 2020

The leave-one-covariate-out conditional randomization test

Eugene Katsevich[1] and Aaditya Ramdas[1,2]

{ekatsevi,aramdas}@stat.cmu.edu
Department of Statistics and Data Science[1]
Machine Learning Department[2]
Carnegie Mellon University

June 16, 2020

Abstract

Conditional independence testing is an important problem, yet provably hard without assumptions. One of the assumptions that has become popular of late is called "model-X", where we assume we know the joint distribution of the covariates, but assume nothing about the conditional distribution of the outcome given the covariates. Knockoffs is a popular methodology associated with this framework, but it suffers from two main drawbacks: only one-bit p-values are available for inference on each variable, and the method is randomized with significant variability across runs in practice. The conditional randomization test (CRT) is thought to be the "right" solution under model-X, but usually viewed as computationally inefficient. This paper proposes a computationally efficient leave-one-covariate-out (LOCO) CRT that addresses both drawbacks of knockoffs. LOCO CRT produces valid p-values that can be used to control the familywise error rate, and has nearly zero algorithmic variability. For L1 regularized M-estimators, we develop an even faster variant called L1ME CRT, which reuses computation by leveraging a novel observation about the stability of the cross-validated lasso to removing inactive variables. Last, for multivariate Gaussian covariates, we present a closed form expression for the LOCO CRT p-value, thus completely eliminating resampling in this important special case.

1 Introduction

1.1 Conditional independence testing under model-X

Testing the importance of variables in high dimensional models is an active area of research with wide-ranging applications. Suppose we have n samples

$$(X_1^i, \dots X_m^i, Y^i) \overset{\text{i.i.d}}{\sim} F_{\boldsymbol{X},\boldsymbol{Y}}, \quad i = 1, \dots, n,$$

from some joint distribution $F_{\boldsymbol{X},\boldsymbol{Y}}$. The null hypothesis that the jth variable is unimportant can be formulated as a statement of conditional independence:

$$H_j : \boldsymbol{Y} \perp\!\!\!\perp \boldsymbol{X}_j \mid \boldsymbol{X}_{-j}. \tag{1}$$

As formalized by Shah and Peters [1], this problem is hard in general: no asymptotically uniformly valid test of this hypothesis can have nontrivial power against any alternatives without assumptions being placed on the distribution $F_{\boldsymbol{X},\boldsymbol{Y}}$. Several works, such as those based on debiasing the lasso [2, 3, 4, 5], have proceeded by placing assumptions like linearity and sparsity on the distribution $F_{\boldsymbol{Y}|\boldsymbol{X}}$. Works on double robustness / double machine learning [6, 7, 8, 1] approach the problem by flexibly learning the distributions $\boldsymbol{Y}|\boldsymbol{X}_{-j}$ and $\boldsymbol{X}_j|\boldsymbol{X}_{-j}$, giving asymptotically valid p-values for certain statistics if at least one of these approximations is accurate enough.

A complementary recent line of work has focused on shifting the assumption to the distribution $F_{\boldsymbol{X}}$. In particular, Candès et al. [9] proposed the

$$\textit{model-X assumption}\text{: the distribution } F_{\boldsymbol{X}} \text{ is known exactly,} \tag{2}$$

while making no assumptions on $F_{\boldsymbol{Y}|\boldsymbol{X}}$. These authors argue that in certain cases, modeling $\boldsymbol{X}$ can be more appropriate than modeling $\boldsymbol{Y}|\boldsymbol{X}$. They also note that a natural way of testing the hypothesis (1) in the model-X framework is the *conditional randomization test* (CRT). The CRT takes an arbitrary test statistic, say the magnitude of a cross-validated lasso coefficient $\widehat{\beta}_j$, and calibrates its distribution under the null by resampling $\boldsymbol{X}_j$ from the conditional distribution $\boldsymbol{X}_j|\boldsymbol{X}_{-j}$. This yields powerful p-values valid in finite samples. However, they point out that the CRT is impractical due to its computational cost. Informally, for sufficient precision to pass a multiple testing correction, we would need each of the p-values to potentially take values close to α/m, and thus we need at least m/α resamples for each of the m variables. Therefore, we need to evaluate the test statistic approximately m^2 times. To get powerful inference, [9] argue that the test statistic should be derived from a multivariate method like the lasso. Therefore, the CRT requires retraining (say) the lasso on m variables m^2 times, which becomes prohibitively expensive for large m.

1.2 Model-X knockoffs: advantages and disadvantages

As an alternative to the CRT, Candès et al. [9] propose *knockoffs*. The idea is to construct knockoff variables $(\widetilde{\boldsymbol{X}}_1, \ldots, \widetilde{\boldsymbol{X}}_m)$ to serve as "negative controls" for the original variables. Then, a multivariate method such as the lasso is run on an augmented design matrix containing the original variables and their knockoffs. Each variable receives a statistic defined as the difference in measured importance between its original and its knockoff. These statistics obey a sign-symmetry property under the null, which can be leveraged to obtain a rigorous false discovery rate (FDR) guarantee. Compared to the CRT, which requires m^2 applications of the lasso involving m variables, knockoffs requires only one application of the lasso involving $2m$ variables. Knockoffs is thus feasible for large-scale problems, and this methodology has successfully been applied to genome-wide association studies (GWAS) [10, 11].

Despite the increasing popularity of knockoffs, this methodology is not without its drawbacks. Principal among these are

- Knockoffs produce only *'one bit p-values'* for each variable.
- The procedure is *randomized*, and different runs produce different outputs.

The term "one bit p-values" refers to the fact that inference is based on the sign of the knockoff statistic, which is positive or negative with equal probability under the null. All the inferential properties of knockoffs are based on this one bit of information per variable. A consequence of this is that no single variable, however strong, can stand on its own; the methodology can only make aggregate statements about multiple variables. Therefore, knockoffs cannot control the family-wise error rate (FWER) at a level of less than 0.5. Controlling the FWER is particularly relevant for GWAS applications, where it is the error rate of choice in the community. Furthermore, practitioners are accustomed to having an interpretable measure of significance, such as a p-value, for each hypothesis. The lack of such a measure for knockoffs makes it hard to interpret the significance of individual variables.

The other drawback of knockoffs is that it is a randomized procedure; one set of knockoff variables is sampled for each observation. This randomization can lead to variability in the performance of the procedure (see Figure 1 below, and Figure 4 of [10]). While it is tempting to generate multiple instances of knockoffs and somehow aggregate the resulting rejection sets ("derandomization"), there is no known principled way of doing this. Incorporating multiple jointly exchangeable knockoffs into the selection procedure itself has been considered [12], though this approach has been found to have low power. In past applications of knockoffs [9, 10], the procedure has been heuristically derandomized by running the entire procedure multiple times and reporting for each variable the fraction of times it ended up in the rejection set. However, the price of this heuristic is a sacrifice of knockoffs' elegant theoretical guarantee and thus interpretability of its results.

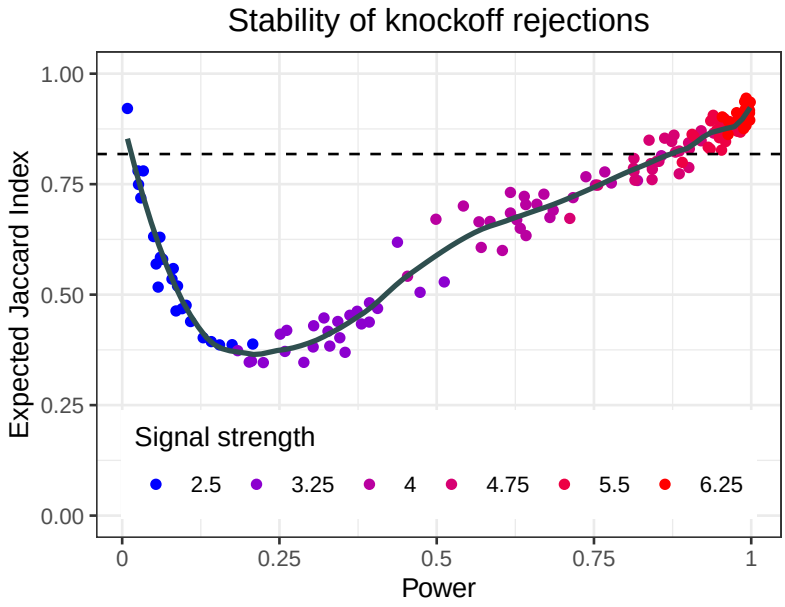

Figure 1: Expected similarity—measured using the Jaccard Index—between knockoff rejection sets resulting from two independent knockoff samples (see definition (17)). Each point corresponds to a realized pair (X, Y), generated from a sparse linear model with 1000 samples, 500 variables, 50 non-nulls, AR(0.5) Gaussian design. There are 25 points (repetitions) each for six signal strengths, and the FDR level is $q = 0.1$ (other details in Section 3). LOESS smoother (gray) is added for visualization. Ideally, the Jaccard index would be around $(1 - q)/(1 + q) \approx 0.82$ (dashed line) if repeated runs lead to only the selected nulls changing with the selected non-nulls staying the same. The dip in the middle of the plot suggests high algorithmic variability in practical settings where the power is not very low or very high.

1.3 Revisiting the conditional randomization test

These drawbacks of knockoffs highlight the appeal of the CRT, which does not suffer from either. This leads to the question: is there a computationally efficient version of the CRT? The recently proposed *holdout randomization test* (HRT) [13] is such a procedure, but it sacrifices power for computation by relying on sample splitting. Indeed, the HRT faces a hard trade-off: each data point must be used either for model fitting or for inference, so power can suffer from either a poor model fit or too few samples for inference. Further, just like knockoffs, the HRT is also randomized, hence both methods suffer from a potential lack of reproducibility. So the HRT only addresses the first drawback above, but not the second.

In this paper, we propose the *leave-one-covariate-out conditional randomization test* (LOCO CRT), a computationally efficient instantiation of the CRT that sacrifices

little or no power compared to the full CRT. Our naming was inspired by the LOCO metric for variable importance suggested recently in the conformal literature [15], but the mathematical guarantees for each are rather different. The LOCO CRT operates by fitting a multivariate regression model to $(\boldsymbol{Y}, \boldsymbol{X}_{-j})$ (i.e. leaving covariate j out) and then regressing the residual on $\boldsymbol{X}_j$. The first step is computationally expensive, but the result can be recycled upon resampling $\boldsymbol{X}_j$. Therefore, each additional resample can be computed in time $O(n)$. Algorithm 1 summarizes the algorithm when the test statistic is an M-estimator. This reduces the number of times we must run a multivariate regression from m^2 to m, a huge gain in high-dimensional settings.

We show that for L_1-regularized M-estimators like the lasso, this computational acceleration can be further improved. We demonstrate that the solution of the lasso with one variable left out can be obtained from the solution of the lasso on the full data, unless the variable is in the active set of the latter. This recycling of computation allows us to compute the m LOCO CRT p-values using only one lasso run on the full data and s runs of the lasso with one variable left out, where s is the size of the full data active set.

Importantly, we show through simulations that unlike HRT, LOCO CRT does not pay much price in power (relative to full CRT) for its computational improvement and is competitive with knockoffs for FDR control in some regimes. Code to reproduce our simulations is available on Github at `https://github.com/ekatsevi/LOCO-CRT`.

Notation. We use bold symbols for population level variables and regular font for their observations in the sample. Let $[a]$ denote the set $\{1, 2, \ldots, a\}$. For consistency, we use $i \in [n]$ to denote an index for data points (always superscript), $j \in [m]$ for variables (always subscript), $g \in [G]$ for gridpoints on a regularization path, and $k \in [K]$ for folds in cross-validation. As an example, $\boldsymbol{X}_j \in \mathbb{R}$ denotes the j-th variables, X_j $(\in \mathbb{R}^n)$ denotes the jth column of the design matrix; i.e. the values of the jth variable for the n samples. The notation X_{-j} $(\in \mathbb{R}^{n \times (m-1)})$ denotes all columns of the design matrix but the jth one. We remark that we use m for number of variables instead of p, because the latter will be used to denote p-values.

Paper outline. In Section 2, we present the proposed LOCO CRT methodology, along with two computational speed-ups. We explore the power of LOCO CRT using numerical simulations in Section 3, and conclude with a discussion in Section 4.

Note: While completing this paper as a methodological counterpart to our theoretical grounding of model-X methods [16], we became aware of a parallel preprint by Liu and Janson [17] with similar goals to this paper. Their ideas appear to have both similarities and differences from ours, but we leave detailed comparisons for later.

2 Leave-one-covariate-out CRT

This section summarizes our three main contributions:

- We propose the LOCO CRT methodology based on any machine learning model for $\boldsymbol{Y}|\boldsymbol{X}_{-j}$. This choice does not affect type-I error but it does affect power, so we choose our algorithm to mimic the optimal test statistic [16].
- We propose the L1ME CRT (pronounced "lime CRT"), an additional speedup to compute the LOCO CRT p-values across all variables when using L1-regularized M-estimators. L1ME CRT operates by recycling computation across variables.
- We derive closed-form p-values and thus avoid resampling altogether when the model for X is multivariate Gaussian, an important special case often considered in the knockoffs literature [18, 19].

2.1 LOCO CRT methodology

For an arbitrary test statistic $T_j(X, Y)$, the CRT p-value [9] is defined

$$p_j(X, Y) \equiv \mathbb{P}[T_j(\widetilde{X}_j, X_{-j}, Y) \geq T_j(X, Y)|X, Y], \tag{3}$$

where $\widetilde{X}_j^i|X, Y \overset{\text{ind}}{\sim} \mathcal{L}(\boldsymbol{X}_j|\boldsymbol{X}_{-j} = X_{-j}^i)$. Since this conditional probability usually cannot be computed explicitly, a p-value is calculated by Monte Carlo based on a large number B of resamples $\widetilde{X}_j^1, \ldots, \widetilde{X}_j^B$:

$$\widehat{p}_j(X, Y) \equiv \frac{1 + \sum_{b=1}^{B} \mathbb{1}(T_j(\widetilde{X}_j^b, X_{-j}, Y) \geq T_j(X, Y))}{B + 1}. \tag{4}$$

We emphasize that p_j and $\widehat{p}_j$ are valid p-values if the model-X assumption holds.

To motivate and derive the LOCO CRT, consider the usual linear model

$$\mathcal{L}(\boldsymbol{Y}|\boldsymbol{X}) = N(\boldsymbol{X}_1\beta_1 + \cdots + \boldsymbol{X}_m\beta_m, \sigma^2), \tag{5}$$

though we remind the reader that no assumption on $F_{\boldsymbol{Y}|\boldsymbol{X}}$ is necessary for Type-I error control. Proposition 1 of our parallel work [16] uses the Neyman-Pearson lemma and Bayes rule to conclude that the optimal CRT statistic is the log-likelihood:

$$T_j^*(X, Y) = -\|Y - X\beta\|^2.$$

Technically, the aforementioned result states that the optimal test statistic is the difference between the log-likelihood of $\boldsymbol{Y}|\boldsymbol{X}$ and $\boldsymbol{Y}|\boldsymbol{X}_{-j}$, which is explicitly written as $-\|Y - X\beta\|^2 + \|Y - X_{-j}\beta_{-j}\|^2$. This form is more interpretable and intuitive,

but as noted there already, the second term behaves like a constant from the point of view of the CRT resampling step, and so the two test statistics will have exactly identical type-I error and power.

Of course, β is unknown, so it is natural to approximate T_j^* by

$$T_j(X, Y) \equiv -\|Y - X\widehat{\beta}\|^2,$$

where $\widehat{\beta}$ is some estimate of the regression coefficients. In a high-dimensional setting, this estimate may be obtained via the lasso. However, as discussed in the introduction and noted earlier [9, 13], re-fitting the lasso model with B resampled design matrices to get the CRT p-value for a single variable is prohibitively expensive. Despite this, one may reasonably assume that most of the computational effort that went into fitting the model once on the original data should somehow be able to be recycled for refitting the model with resampled design matrices (since only column j is resampled in each step of the CRT, and thus most of the data matrix remains the same).

Indeed, there is an elegant way of recycling computation. Our starting point is a fact from low-dimensional linear regression theory: the ordinary least squares coefficient $\widehat{\beta}_j$ can be obtained by first regressing X_{-j} out of Y and separately out of X_j, and then regressing the residual from Y on the residual from X_j. This motivates the following procedure to obtain a test statistic:

1. Use any method to learn $\widehat{\beta}_{-j} \in \mathbb{R}^{m-1}$ based on (Y, X_{-j}), and extract the residual $r_j \equiv Y - X_{-j}\widehat{\beta}_{-j}$.
2. Compute $\mu_j \equiv \mathbb{E}[X_j|X_{-j}]$ based on the model-X assumption (2).
3. Regress r_j on $X_j - \mu_j$ to obtain $\widehat{\beta}_j \equiv \dfrac{(X_j - \mu_j)^T r_j}{\|X_j - \mu_j\|^2}$.
4. Compute the test statistic $T_j(X, Y) \equiv -\|Y - X_{-j}\widehat{\beta}_{-j} - (X_j - \mu_j)\widehat{\beta}_j\|^2$.

Now, note that steps 1 and 2 involve functions of only (Y, X_{-j}), so only steps 3 and 4 must be repeated when X_j is replaced by $\widetilde{X}_j$. Therefore, each resampling step does not require a new model fit and takes time only $O(n)$ (if sampling $\widetilde{X}_j$ is fast). It usually takes less time to do all the resampling than to fit the original model.

These computational savings come at negligible statistical price. Indeed, if we consider the OLS example, the above test statistic is essentially the same as directly fitting $\widehat{\beta}$. If $\widehat{\beta}_{-j}$ is fit using the lasso, the coefficient vector $\widehat{\beta}$ obtained from the above steps will not be the same as that obtained from running a lasso on all the variables. Nevertheless, this change will not have a large impact on the power of the CRT; we may as well explain as much of variance in Y as possible using X_{-j} because the power to reject (1) lies only in the part of Y not explained by X_{-j}.

The above intuition leads to Algorithm 1, a general definition of the proposed LOCO CRT procedure. We replace the squared error with an arbitrary differentiable and strictly convex loss function ℓ, and we do not assume linearity as in (5).

Algorithm 1: LOCO CRT for variable j

Data: Pairs (X^i, Y^i) for $i = 1, \ldots, n$, loss function ℓ, machine learning method f

Fit f to (Y, X_{-j}) to obtain $\widehat{f}_{-j}$, and calculate fitted values $o_i \equiv \widehat{f}_{-j}(X^i_{-j})$;

Compute $\mu_j \equiv \mathbb{E}[X_j | X_{-j}]$;

Solve univariate problem with offsets: $\widehat{\beta}_j \equiv \min_{\beta_j} \sum_{i=1}^n \ell(Y^i, (X^i_j - \mu_j)\beta_j + o_i)$;

Compute $T_j(X, Y) \equiv \sum_{i=1}^n \ell(Y^i, (X^i_j - \mu_j)\widehat{\beta}_j + o_i)$;

for $b \in \{1, \ldots, B\}$ **do**

  Sample $\widetilde{X}^b_j$ from $\mathcal{L}(X_j | X_{-j})$;

  Compute $T_j(\widetilde{X}^b_j, X_{-j}, Y)$ by repeating steps 3 and 4 with $\widetilde{X}^b_j$ replacing X_j;

end

Compute $\widehat{p}_j$ based on formula (4);

Result: LOCO CRT Monte Carlo p-value $\widehat{p}_j$.

In practice, Algorithm 1 would be applied for each variable j to obtain CRT p-values $p_1, \ldots, p_m$, where p_j tests the hypothesis (1). The computational cost of Algorithm 1 is $m \cdot (T_f + B \cdot n)$, where T_f is the cost of running the machine learning algorithm to fit $\widehat{f}^{(j)}_{-j}$. In comparison, the cost of running the full CRT (without LOCO) is $m \cdot B \cdot T_f$. Therefore, the computational cost of LOCO CRT is B times smaller than full CRT, and this difference can be several orders of magnitude. The cost of running the HRT, $T_f + m \cdot Bn$, is even smaller, but at the cost of lower power and higher algorithmic variability. By contrast, the computational cost of knockoffs is approximately T_f, as the machine learning algorithm needs to be fit only once. Table 1 summarizes the above discussion.

Model-X method	Total computation time	Algorithmic variability
Knockoffs [9]	$T_{f,2m}$	Non-negligible
Full CRT [9]	$m \cdot B \cdot T_{f,m}$	Negligible
HRT [13]	$T_{f,m} + m \cdot Bn$	Non-negligible
LOCO CRT (Algorithm 1)	$m \cdot (T_{f,m-1} + Bn)$	Negligible
L1ME CRT (Algorithm 2)	$\lvert\mathcal{A}\rvert \cdot (T_{f,m-1} + Bn)$	Negligible

Table 1: Summary of the computation times and algorithmic variabilities of model-X methods based on n samples. $T_{f,m}$ is the time to fit a machine learning algorithm on m variables. $\mathcal{A}$ is the lasso active set, B is the number of Monte Carlo resamples.

We end by noting that the HRT splits across data points, but LOCO CRT splits across variables (it leaves one variable out, as opposed to one data point out). While both strategies save computation, they are crucially different. The HRT requires just one model fit, but sacrifices statistical efficiency by splitting the data. The LOCO CRT, though much faster than full CRT, still requires fitting a model for each variable. But as we see in Section 3, LOCO CRT is as powerful as full CRT and much more powerful than HRT. A natural question is whether the strengths of the two strategies can be combined, but we leave this for future work.

Despite the fact that the LOCO CRT is fully parallelizable across variables, the cost $m \cdot T_f$ might still be prohibitive for massive datasets. In the next subsection, we propose a computational shortcut for lasso-based test statistics that computes all LOCO CRT p-values in time $|\mathcal{A}| \cdot T_f$, where $\mathcal{A}$ is the active set of the cross-validated lasso applied to all variables.

2.2 L1ME CRT for L1-regularized M-estimators

Let us consider the case of a general L1-regularized M-estimator:

$$\widehat{\beta}(X, Y; \lambda) \equiv \underset{\beta \in \mathbb{R}^m}{\arg\min} \sum_{i=1}^{n} \ell(Y^i, X^i\beta) + \lambda \|\beta\|_1 , \tag{6}$$

where $\lambda > 0$. To avoid degeneracies, we assume throughout that the columns of X are in general position. This, together with the strict convexity and differentiability of ℓ, guarantees that $\widehat{\beta}$ is unique [20]. Suppose we wish to apply Algorithm 1 for all variables j, and would like to use the lasso to fit each $\widehat{f}_{-j} \equiv X_{-j}\widehat{\beta}(X_{-j}, Y; \lambda)$, where

$$\widehat{\beta}(X_{-j}, Y; \lambda) \equiv \underset{\beta \in \mathbb{R}^{m-1}}{\arg\min} \sum_{i=1}^{n} \ell(Y^i, X^i_{-j}\beta) + \lambda \|\beta\|_1 . \tag{7}$$

The LOCO problem (7) is quite similar to the full problem (6) for each j. In fact, for a fixed regularization parameter λ, we show that the solution for the former can often be directly derived from the solution for the latter.

Lemma 2.1. *Suppose the columns of X are in general position and that the loss ℓ is differentiable and strictly convex. If $\widehat{\beta}_j(X, Y; \lambda) = 0$ for some λ and some variable j, then the LOCO lasso coefficient vector is the same as the full lasso coefficient vector with the jth coordinate removed:*

$$\widehat{\beta}_j(X, Y; \lambda) = 0 \quad \Longrightarrow \quad \widehat{\beta}(X_{-j}, Y; \lambda) = \widehat{\beta}_{-j}(X, Y; \lambda). \tag{8}$$

The above lemma (proved in the appendix) has important computational implications, because it suggests that we can avoid refitting the LOCO lasso for most

variables j, instead recycling the lasso fit on the full design matrix. Of course, the parameter λ is usually tuned via cross-validation, which introduces extra complications. However, we claim that if λ is chosen in an appropriate data-dependent way, then the result (8) will still hold for most j.

To make this precise, first let $\{1, \ldots, n\} = D_1 \cup \cdots \cup D_K$ be a partition of the data points into K folds and let

$$\lambda(1) > \lambda(2) > \cdots > \lambda(G) > 0 \tag{9}$$

be a grid of regularization parameters. Now, consider applying K-fold cross-validation to (X_{-j}, Y). For each fold k and each grid point g, we compute $\widehat{\beta}(X_{-j}^{-D_k}, Y^{-D_k}; \lambda(g))$, the coefficient estimate on the data obtained by holding out the kth fold and the jth variable. Summing over variables, define the CV error for the jth variable when using λ_g as

$$\mathcal{E}_{-j,g} \equiv \sum_{i=1}^{n} \ell(Y^i, X_{-j}^i \widehat{\beta}(X_{-j}^{-D_k}, Y^{-D_k}; \lambda(g)),$$

where $k(i)$ is the fold k to which data point i belongs. Now, we select $\widehat{\lambda} = \lambda(\widehat{g}_{-j})$ via some rule $\widehat{g}$ as follows:

$$\widehat{g}_{-j} \equiv \widehat{g}(\mathcal{E}_{-j,1}, \ldots, \mathcal{E}_{-j,G}), \tag{10}$$

and then re-run the lasso on all data points to get the final estimate $\widehat{\beta}(X_{-j}, Y; \lambda(\widehat{g}_{-j}))$.

In the spirit of Lemma 2.1, we would like to run this procedure on the full data, and then recycle the results to recover $\widehat{\beta}(X_{-j}, Y; \lambda(\widehat{g}_{-j}))$ for most j. To obtain such a property, we must restrict our attention to *sequential* rules to select $\widehat{\lambda}$:

Definition 1. *A rule $\widehat{g}$ to select the penalty parameter λ based on CV errors $\mathcal{E}_1, \ldots, \mathcal{E}_G$ is "sequential" if these values are traversed in this order, and at some stopping time $\widetilde{g}$, the algorithm terminates and chooses $\lambda(\widehat{g})$ for some $\widehat{g} \leq \widetilde{g}$.*

For example, for any integer $\Delta \geq 1$, the following rule is sequential:

$$\widehat{g} \equiv \min\{g : \mathcal{E}_g \leq \min(\mathcal{E}_{g+1}, \ldots, \mathcal{E}_{g+\Delta})\},$$

which is the first time along the regularization path that the CV error is smaller than the following Δ steps (the first 'local minimum' on the CV path, and the 'sparsest' of all such local minima). In this case the stopping time is $\widetilde{g} = \widehat{g} + \Delta$.

The following novel proposition gives a key computational speedup; its proof can be found in Appendix A.

Proposition 2.2. *Fix a grid of regularization parameters (9). Consider applying the ℓ_1-regularized regression (6) on the whole data (X, Y), with λ selected by cross-validation with a sequential stopping rule $\widehat{g}$. Let $\widehat{g}_{\text{full}}$ and $\widetilde{g}_{\text{full}}$ be the resulting grid point and stopping time, respectively. Define the active set*

$$\begin{aligned}\mathcal{A} \equiv \{j \in [m] : \widehat{\beta}_j(X, Y; \lambda(\widehat{g}_{\text{full}})) \neq 0 \text{ or} \\ \widehat{\beta}_j(X^{-D_k}, Y^{-D_k}; \lambda(g)) \neq 0 \text{ for some } k, g \leq \widetilde{g}_{\text{full}}\}.\end{aligned} \tag{11}$$

If the loss ℓ is differentiable and strictly convex, and the columns of X and X^{-D_k} are in general position for each k, then excluding non-active variables j does not alter the fitted coefficients:

$$\text{for each } j \notin \mathcal{A}, \quad \widehat{g}_{-j} = \widehat{g}_{\text{full}} \quad \text{and} \quad \widehat{\beta}(X_{-j}, Y; \lambda(\widehat{g}_{-j})) = \widehat{\beta}_{-j}(X, Y; \lambda(\widehat{g}_{\text{full}})). \tag{12}$$

Proposition 2.2 states that for each variable j *not* in the active set, we need not re-run the lasso holding out variable j; we can instead fit the full lasso once and then read off the coefficient vector. This computational shortcut, summarized in Algorithm 2, reduces the computational cost of the LOCO CRT from $m \cdot C$ to $|\mathcal{A}| \cdot C$. Depending on the sparsity of the problem, this reduction can save several orders of magnitude of computation. It is known that at most, the lasso solution has $\min(m, n)$ nonzero entries [20], though in most cases it is much sparser.

Algorithm 2: L1ME CRT

Data: Pairs (X^i, Y^i) for $i = 1, \ldots, n$, sequence $\lambda(1) > \cdots > \lambda(G) > 0$, loss ℓ, sequential rule $\widehat{g}$

Fit a cross-validated lasso on the full design matrix (6); choose $\widehat{g}_{\text{full}}$ via the rule $\widehat{g}$ to obtain $\widehat{\beta}(X, Y; \lambda(\widehat{g}_{\text{full}}))$; record the active set $\mathcal{A}$ as defined in (11);

for $j \in \mathcal{A}$ **do**

  Refit the lasso on the LOCO design X_{-j} using the rule $\widehat{g}$ to obtain $\widehat{\beta}_{-j} \equiv \widehat{\beta}(X_{-j}, Y; \lambda(\widehat{g}_{-j}))$.

end

for $j \notin \mathcal{A}$ **do**

  Set $\widehat{\beta}_{-j} = \widehat{\beta}_{-j}(X, Y; \lambda(\widehat{g}_{\text{full}}))$.

end

For each j, let $\widehat{f}_{-j}(X_{-j}) = X_{-j}\widehat{\beta}_{-j}$;

Result: Functions $\widehat{f}_{-j}$ to be used with Algorithm 1 for all j.

2.3 Exact LOCO CRT for Gaussian covariates

In line with the literature on knockoffs, we consider the special case of model-X when the covariates are normally distributed:

$$\boldsymbol{X} \sim N(0, \Sigma) \tag{13}$$

for some covariance matrix $\Sigma \in \mathbb{R}^{m\times m}$. In this case, we show that LOCO CRT and HRT p-values (3) can be computed exactly without resorting to Monte Carlo. Proofs are deferred to Appendix A. Consider the following statistic using the squared loss:

$$T_j(X,Y) \equiv -\|Y - \widehat{f}_{-j}(X_{-j}) - (X_j - \mu_j)\widehat{\beta}_j\|^2 \equiv -\|r_j - (X_j - \mu_j)\widehat{\beta}_j\|^2. \tag{14}$$

For LOCO CRT, we define

$$\widehat{\beta}_j \equiv \frac{(X_j - \mu_j)^T r_j}{\|X_j - \mu_j\|^2},$$

and for HRT, we learn $\widehat{\beta}_j$ based on the training set. In both of these cases, there is an exact expression for the p-value based on statistic (14).

Proposition 2.3. *Suppose the covariates have a Gaussian distribution (13) and ℓ is squared error loss. Then, the LOCO CRT p-value has the following explicit form:*

$$p_j^{LOCO}(X,Y) = \mathbb{P}\left[F_{n-1,1} \le \frac{1}{n-1}\left(\left(\frac{(X_j - \mu_j)^T(Y - \widehat{f}_{-j}(X_{-j}))}{\|X_j - \mu_j\| \cdot \|Y - \widehat{f}_{-j}(X_{-j})\|}\right)^{-2} - 1\right)\middle| X, Y\right], \tag{15}$$

where $F_{n-1,1}$ represents a random variable with an F distribution on $n-1$ and 1 degrees of freedom.

In the Gaussian case, the LOCO CRT p-value is therefore exactly the tail probability of an F distribution. The formula (15) also shows that the test statistic is equivalent to the empirical correlation between X_j and Y, after both are corrected for X_{-j}.

Proposition 2.4. *Let $n = n_{\text{train}} + n_{\text{test}}$ be a split of the samples for training and testing. Suppose the covariates have a Gaussian distribution (13). Consider the HRT defined by a model $\widehat{f}(X) = \widehat{f}_{-j}(X_{-j}) + (X_j - \mu_j)\widehat{\beta}_j$, trained on n_{train} data points, with squared error loss. Then, the HRT p-value has the following explicit form:*

$$p_j^{HRT}(X,Y) = \mathbb{P}\left[\chi^2_{n_{\text{test}},\text{ncp}} \le \frac{\|Y - \widehat{f}(X)\|^2}{s_j^2\widehat{\beta}_j^2}\middle| X, Y\right]; \quad \text{ncp} \equiv \frac{\|Y - \widehat{f}_{-j}(X_{-j})\|^2}{s_j^2\widehat{\beta}_j^2}, \tag{16}$$

where $\chi^2_{d,\text{ncp}}$ denotes a noncentral chi-square random variable with d degrees of freedom and noncentrality parameter ncp *and $s_j^2 \equiv \text{Var}[\boldsymbol{X}_j|\boldsymbol{X}_{-j}]$.*

Thus, the HRT p-value also has an explicit expression, though it is the tail probability of a noncentral chi-square distribution rather than an F distribution. This difference arises from the fact that HRT fits $\widehat{\beta}_j$ out-of-sample, while the LOCO CRT fits it in-sample.

3 Numerical simulations

In this section, we evaluate the performance of the LOCO CRT on simulated data; code to reproduce our experiments is available at https://github.com/ekatsevi/LOCO-CRT. All experiments are based on a sparse linear model (5) with Gaussian covariates (13). We use an autocorrelated design matrix $\Sigma_{j_1,j_2} = \rho^{|j_1-j_2|}$, with $\rho \in \{0.2, 0.5, 0.8\}$. The k non-null entries of the coefficient vector β are equally spaced, all have magnitude $A/\sqrt{n}$, and random signs. We call A the *signal strength*.

LOCO CRT has similar power to full CRT. Here, we take $n = 200$ and $m = 1000$, with $k = 50$ non-null coefficients. We assess the power of following four methods to reject a single non-null variable at level $\alpha = 0.2$:

- Full CRT: Using the 10-fold cross-validated lasso coefficient, leaving the variable of interest unpenalized.
- LOCO CRT: Algorithm 1 with squared loss, accelerated using Proposition 2.3.
- HRT: Learn a linear model using 10-fold cross-validated lasso on 50% of the data (variable of interest unpenalized), and then test on the remaining 50% using the squared error loss. Accelerated using Proposition 2.4.
- Oracle HRT: Same as the HRT, but using the ground truth model instead of the learned model.

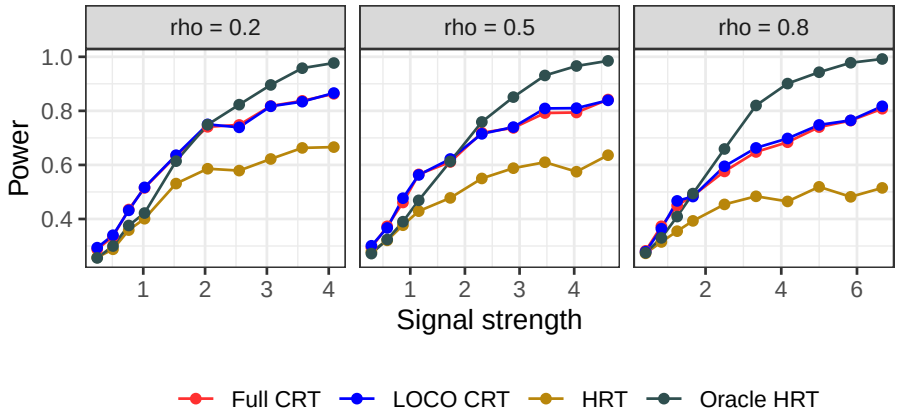

Figure 2: Power of CRT and HRT variants to reject one non-null variable at $\alpha = 0.2$. Each point computed based on 1000 repetitions. LOCO CRT performs similarly to full CRT and outperforms HRT.

Figure 2 shows the power of each method versus the signal strength (a normalized version of the coefficient magnitudes) across the three correlation settings. We see that the power of LOCO CRT tracks that of the full CRT quite closely, suggesting that the computational advantage of LOCO comes without much of a statistical price. On the other hand, the HRT performs significantly worse than the two CRT variants, demonstrating that sample splitting does come with a statistical cost. Finally, even the oracle HRT methodology can sometimes perform worse than the (LOCO) CRT—in spite of the fact that it uses the ground truth model—due to the smaller number of samples it has access to for inference.

LOCO CRT controls FWER more powerfully than HRT. Next, we assess the FWER control performance of the following three methods in a problem with $m = 500$ and $n = 1000$:

- The L1ME CRT (Algorithm 2), accelerated using Proposition 2.3.
- The HRT with 80%/20% train/test split, accelerated using Proposition 2.4.
- The CRT based on marginal correlation $T_j(X, Y) \equiv (X_j - \mu_j)^T Y$.

We apply each method to all 500 variables, which excludes the full CRT due to its computational complexity. We Bonferroni adjust the p-values returned by each method with the goal of controlling the FWER at level $\alpha = 0.05$. We repeat the experiment 25 times each for ten signal strengths, $\rho \in \{0.2, 0.5, 0.8\}$, and $k \in \{25, 50, 100\}$.

As discussed in the introduction, knockoffs cannot control the FWER at this level due to its reliance on one-bit p-values. Figure 3 compares the power of these three methods, and it shows that the LOCO CRT significantly outperforms the HRT. It also outperforms the CRT based on marginal correlation, where the gap depends on how well LOCO CRT learns the contributions of the other variables to the response.

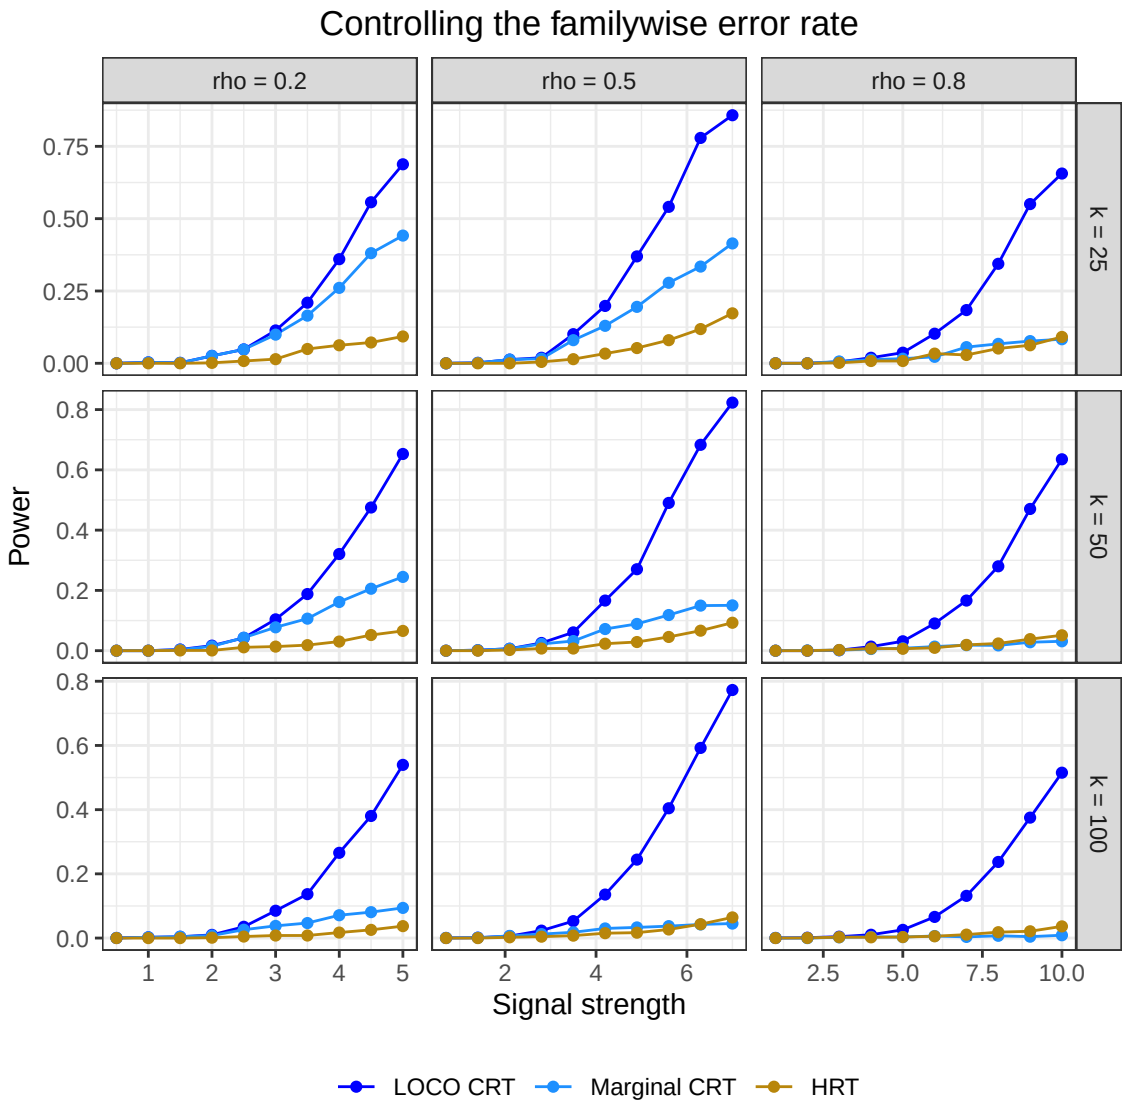

Figure 3: Comparing the power of three CRT/HRT variants controlling the FWER at level $\alpha = 0.05$. LOCO CRT has better power than the two available alternatives (notably missing knockoffs), because it leverages the power of multivariate modeling without the need to split the sample.

Model-X knockoffs can outperform LOCO CRT when controlling FDR, but exhibit substantial algorithmic variability. Next, we applied model-X knockoffs, in addition to the three above CRT/HRT variants with the BH procedure, to control the FDR at level $q = 0.1$ in the same problem setting. Figure 4 shows the power of these four methods. LOCO CRT continues to outperform HRT and marginal CRT. It also performs as well as or better than knockoffs in the case $\rho = 0.2$. However, as the correlation increases, knockoffs surpasses the LOCO CRT. This discrepancy is intriguing and deserves further attention.

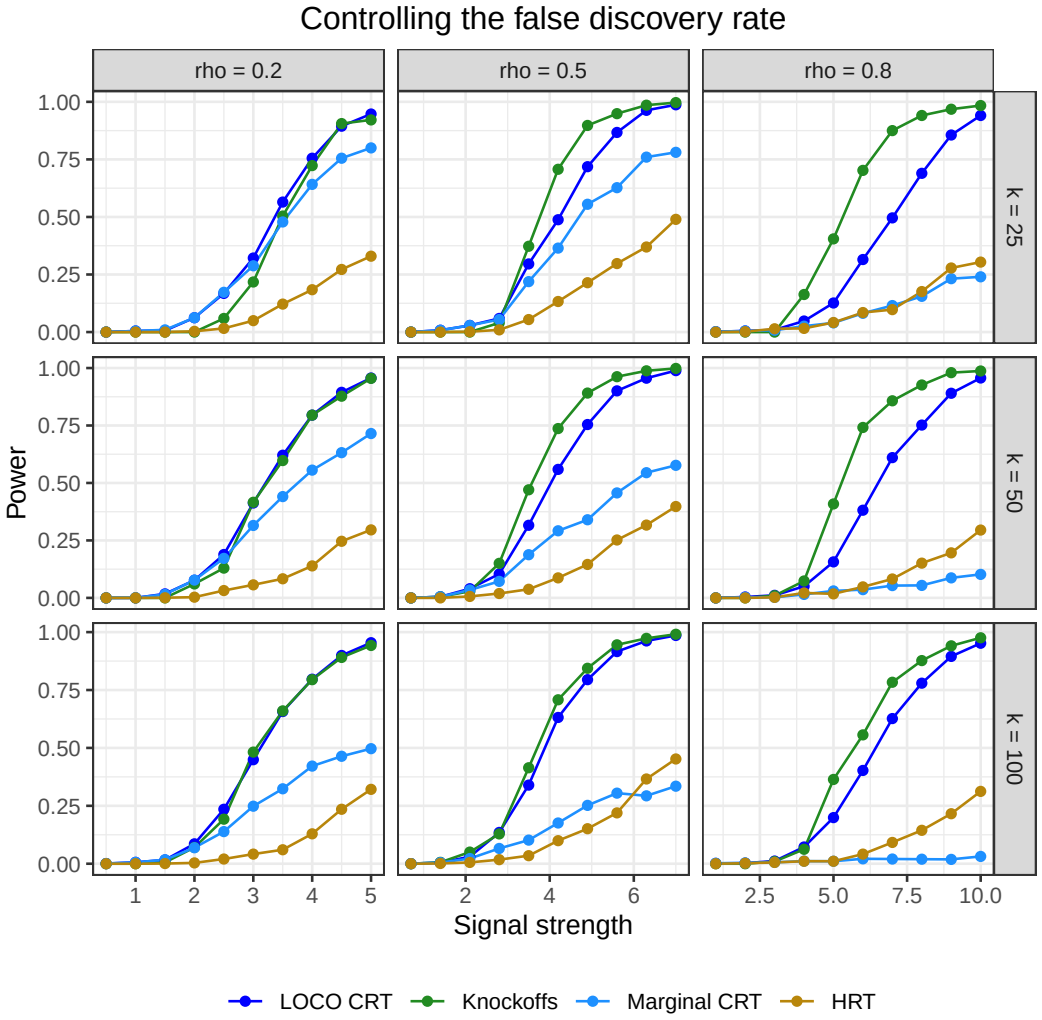

Figure 4: Comparing the power three methods from Figure 3 to that of knockoffs, this time controlling the FDR at level $q = 0.1$. LOCO CRT outperforms or matches knockoffs for low correlation, but has lower power for medium and high correlations.

Despite the good average power of knockoffs, we find that it suffers generally poor stability with respect to knockoff resampling. Given data X, Y, consider two independent draws of the knockoff matrix $\widetilde{X}^{(1)}$ and $\widetilde{X}^{(2)}$. These lead to two rejection sets $\mathcal{R}^{(r)} \equiv \mathcal{R}([X, \widetilde{X}^{(r)}], Y)$, $r = 1, 2$. Their similarity can be measured via the *Jaccard Index*, defined

$$J(\mathcal{R}^{(1)}, \mathcal{R}^{(2)}) \equiv \frac{|\mathcal{R}^{(1)} \cap \mathcal{R}^{(2)}|}{|\mathcal{R}^{(1)} \cup \mathcal{R}^{(2)}|}.$$

This quantity is between 0 and 1, with higher Jaccard Indices representing greater similarity. We define the *stability* of the knockoffs rejection set $\mathcal{R}([X, \widetilde{X}], Y)$ as the expectation of this quantity:

$$\text{stability}(X, Y) \equiv \mathbb{E}\left[J\left(\mathcal{R}([X, \widetilde{X}^{(1)}], Y), \mathcal{R}([X, \widetilde{X}^{(2)}], Y)\right) \middle| X, Y \right]. \tag{17}$$

To assess this measure of stability in the above simulation, we reran each experiment with 50 independently drawn knockoffs realizations. Then, we approximated the stability via

$$\widehat{\text{stability}}(X, Y) \equiv \frac{1}{\binom{50}{2}} \sum_{1 \le r_1 < r_2 \le 50} J\left(\mathcal{R}([X, \widetilde{X}^{(r_1)}], Y), \mathcal{R}([X, \widetilde{X}^{(r_2)}], Y)\right).$$

Figure 5, whose middle panel was presented in Figure 1, shows the stability of knockoffs rejections across different values of ρ and k. If only the null rejections changed with knockoff resampling, we would expect the stability to be $(1-q)/(1+q) \approx 0.82$, which is denoted by the dashed line in the figure. Aside from the extreme scenarios when the power is either very low or very high, the stability of the knockoffs rejection set generally fails to meet this threshold. This suggests that even the non-nulls rejected by knockoffs change based on the randomization.

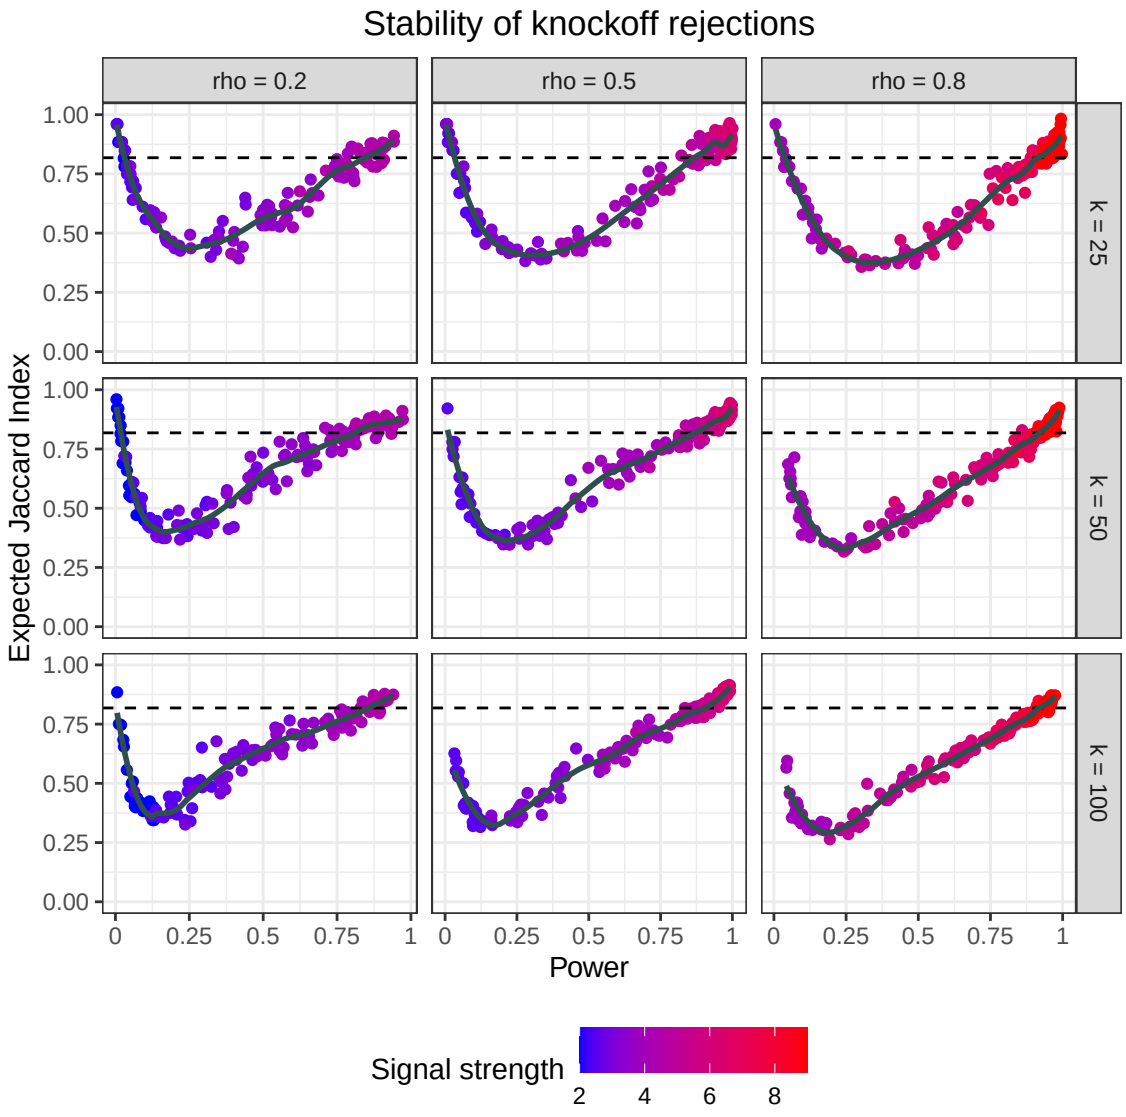

Figure 5: Knockoffs variability as measured by expected Jaccard Index between two independent knockoff realizations, plotted against realized power. Details are as in Figure 1, which displayed the middle panel. This significant variability across re-runs is acknowledged to be a drawback of knockoffs from a reproducibility standpoint.

4 Discussion

The HRT provided the first indication that a variant of the CRT could be computationally tractable, albeit at the cost of statistical performance. In this paper, we demonstrate that leaving out *variables* instead of *samples* creates a procedure that is almost as fast but much more powerful. While not yet quite as powerful or computationally efficient as knockoffs, the negligible algorithmic variability, easy-to-understand procedure, and fine-grained p-value output makes the LOCO CRT an attractive alternative.

We contend that returning a p-value for the significance of each variable is a particularly important advantage of CRT-style methods. Beyond their familiarity and interpretability, p-values allow for various downstream analyses. For example, they open the door to the application of decades worth of multiple testing methodology to the variable selection problem. Knockoffs methodology, though it does not yield p-values, has nevertheless been extended to carry out an impressive variety of multiple testing tasks (Table 2). Other tasks still remain out of reach for knockoffs. With the advent of fast CRT methods like LOCO CRT, these statistical tasks may not need to be challenges requiring custom novel solutions, instead becoming opportunities to employ existing multiple testing methodology in the model-X framework.

Multiple testing task	Approach based on p-values	Knockoffs approach
Global null testing	Fisher '25 [21]	Duan et al. '19 [22]
FWER control	Bonferroni '36 [23]	[Impossible]
Null proportion estimation	Schweder, Spjøtvoll '82 [24]	Katsevich, Ramdas '20[1] [25]
FDR control	Benjamini, Hochberg '95 [26]	Barber, Candès '15 [18]
k-FWER control	Lehmann, Romano '05 [27]	Janson, Su '16 [28]
FDR control with side info	Genovese et al. '06 [29]	Lei, Fithian '19 [30, 31, 32]
Hierarchical FWER control	Meinshausen '08 [33, 34, 35]	[Impossible]
Hierarchical FDR control	Yekutieli '08 [36, 37, 38, 39]	Lei et al. '20 [31]
Group FDR control	Hu et al. '10 [40]	Dai, Barber '16 [41]
Simultaneous FDP control	Goeman, Solari '11 [42]	Katsevich, Ramdas '20 [25]
Multilayer FDR control	Barber, Ramdas '15 [43]	Katsevich, Sabatti '19 [44]

Table 2: Multiple testing tasks, p-value based methodologies to address them, and their knockoff counterparts if applicable. The authors and dates listed correspond to the earliest references we are aware of.

[1]Though Katsevich and Ramdas [25] do not discuss null proportion estimation directly, it has since been pointed out [45] that simultaneous FDP controlling procedures can be applied to null proportion estimation via closed testing.

There is still much to be done on this front. We may not always be able to plug-and-play CRT p-values in multiple testing procedures, since their dependency structure is currently unknown. Nevertheless, we note that Bates et al. [46] recently proposed a clever method of generating independent HRT p-values for groups of linearly-structured covariates, at a modest cost in power. Furthermore, many multiple testing procedures are fairly robust or entirely agnostic to p-value dependence.

The LOCO CRT is a proof of principle, but can likely be fine-tuned for better power and even faster computation. Many degrees of freedom in the construction of LOCO CRT test statistic remain to be explored. For instance, should the statistic be based on the fitted coefficient of a variable or on the loss function? If the latter, should the loss be evaluated in-sample (as is done here) or out-of-fold? The theory we recently developed [16] may help guide the search for powerful test statistics. Finally, we are currently attempting to extend the exact expressions derived in Section 2.3 to discrete covariates.

We are optimistic about the prospects of the CRT for fruitful practical application, and look forward to continued improvements in the computational and statistical efficiency of model-X methodology.

Acknowledgments

This work used the Extreme Science and Engineering Discovery Environment (XSEDE) [47], which is supported by National Science Foundation grant number ACI-1548562. Specifically, it used the Bridges system [48], which is supported by NSF award number ACI-1445606, at the Pittsburgh Supercomputing Center (PSC).

References

[1] Rajen D. Shah and Jonas Peters. The Hardness of Conditional Independence Testing and the Generalised Covariance Measure. *Annals of Statistics, to appear*, 2020.

[2] Adel Javanmard and Andrea Montanari. Confidence intervals and hypothesis testing for high-dimensional regression. *Journal of Machine Learning Research*, 15(1):2869–2909, 2014.

[3] Adel Javanmard and Andrea Montanari. Hypothesis testing in high-dimensional regression under the gaussian random design model: Asymptotic theory. *IEEE Transactions on Information Theory*, 60(10):6522–6554, 2014.

[4] Cun-Hui Zhang and Stephanie S Zhang. Confidence intervals for low dimensional parameters in high dimensional linear models. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 76(1):217–242, 2014.

[5] Sara Van De Geer, Peter Bühlmann, Ya'acov Ritov, and Ruben Dezeure. On asymptotically optimal confidence regions and tests for high-dimensional models. *Annals of Statistics*, 42(3):1166–1202, 2014.

[6] James M. Robins and Andrea Rotnitzky. Comment on the Bickel and Kwon article,Inference for semiparametric models: Some questions and an answer. *Statistica Sinica*, 11(4):920–936, 2001.

[7] Mark J. van der Laan and James M. Robins. *Unified methods for censored longitudinal data and causality*. Springer-Verlag, New York, 2003.

[8] Victor Chernozhukov, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins. Double/debiased machine learning for treatment and structural parameters. *Econometrics Journal*, 21(1):C1–C68, 2018.

[9] Emmanuel Candès, Yingying Fan, Lucas Janson, and Jinchi Lv. Panning for gold: 'model-X' knockoffs for high dimensional controlled variable selection. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(3):551–577, 2018.

[10] M. Sesia, C. Sabatti, and E. J. Candès. Gene hunting with hidden Markov model knockoffs. *Biometrika*, 106(1):1–18, 2019.

[11] Matteo Sesia, Eugene Katsevich, Stephen Bates, Emmanuel Candès, and Chiara Sabatti. Multi-resolution localization of causal variants across the genome. *Nature Communications*, 11:1093, 2020.

[12] Jaime Roquero Gimenez and James Zou. Improving the Stability of the Knockoff Procedure: Multiple Simultaneous Knockoffs and Entropy Maximization. *Proceedings of Machine Learning Research*, 89:2184–2192, 2019.

[13] Wesley Tansey, Victor Veitch, Haoran Zhang, Raul Rabadan, and David M. Blei. The Holdout Randomization Test: Principled and Easy Black Box Feature Selection. *arXiv*, 2018.

[14] Ruben Dezeure, Peter Bühlmann, Lukas Meier, and Nicolai Meinshausen. High-dimensional inference: Confidence intervals, P-values and R-software hdi. *Statistical Science*, 30(4):533–558, 2015.

[15] Jing Lei, Max G'Sell, Alessandro Rinaldo, Ryan J. Tibshirani, and Larry Wasserman. Distribution-Free Predictive Inference for Regression. *Journal of the American Statistical Association*, 113(523):1094–1111, 2018.

[16] Eugene Katsevich and Aaditya Ramdas. A theoretical treatment of conditional independence testing under Model-X. *arXiv*, 2020.

[17] Molei Liu and Lucas Janson. Fast and Powerful Conditional Randomization Testing via Distillation. *arXiv*, 2020.

[18] Rina Foygel Barber and Emmanuel J Candès. Controlling the false discovery rate via knockoffs. *The Annals of Statistics*, 43(5):2055–2085, 2015.

[19] Dongming Huang and Lucas Janson. Relaxing the Assumptions of Knockoffs by Conditioning. *Annals of Statistics, to appear*, 2020.

[20] Ryan J Tibshirani. The lasso problem and uniqueness. *Electronic Journal of Statistics*, 7:1456–1490, 2013.

[21] R A Fisher. *Statistical methods for research workers*. Oliver and Boyd, Edinburgh, 1925.

[22] Boyan Duan, Aaditya Ramdas, Sivaraman Balakrishnan, and Larry Wasserman. Interactive Martingale Tests for the Global Null. *arXiv*, 2019.

[23] C Bonferroni. Teoria statistica delle classi e calcolo delle probabilita. *Pubblicazioni del R Istituto Superiore di Scienze Economiche e Commericiali di Firenze*, 8:3–62, 1936.

[24] T. Schweder and E. Spjotvoll. Plots of P-Values to Evaluate Many Tests Simultaneously. *Biometrika*, 69(3):493–502, 1982.

[25] Eugene Katsevich and Aaditya Ramdas. Simultaneous high-probability bounds on the false discovery proportion in structured, regression, and online settings. *The Annals of Statistics, to appear*, 2020.

[26] Yoav Benjamini and Yosef Hochberg. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 57(1):289–300, 1995.

[27] E L Lehmann and J Romano. Generalizations of the familywise error rate. *The Annals of Statistics*, 33(3):1138–1154, 2005.

[28] Lucas Janson and Weijie Su. Familywise error rate control via knockoffs. *Electronic Journal of Statistics*, 10:960–975, 2016.

[29] C Genovese, K Roeder, and L Wasserman. False discovery control with p-value weighting. *Biometrika*, 93(3):509–524, 2006.

[30] Lihua Lei and William Fithian. AdaPT: an interactive procedure for multiple testing with side information. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 80(4):649–679, 2018.

[31] Lihua Lei, Aaditya Ramdas, and William Fithian. STAR: A general interactive framework for FDR control under structural constraints. *Biometrika, to appear*, 2020.

[32] Zhimei Ren and Emmanuel Candès. Knockoffs with Side Information. *arXiv*, pages 1–29, 2020.

[33] Nicolai Meinshausen. Hierarchical testing of variable importance. *Biometrika*, 95(2):265–278, 2008.

[34] Rosa J. Meijer and Jelle J. Goeman. A multiple testing method for hypotheses structured in a directed acyclic graph. *Biometrical Journal*, 57(1):123–143, 2015.

[35] Rosa J. Meijer and Jelle J. Goeman. Multiple testing of gene sets from Gene Ontology: possibilities and pitfalls. *Briefings in Bioinformatics*, 17(5):808–818, 2015.

[36] Daniel Yekutieli. Hierarchical false discovery rate-controlling methodology. *Journal of the American Statistical Association*, 103(481):309–316, 2008.

[37] Gavin Lynch. *The Control of the False Discovery Rate Under Structured Hypotheses*. PhD thesis, New Jersey Institute of Technology, Department of Mathematical Sciences, 2014.

[38] Aaditya Ramdas, Jianbo Chen, Martin J. Wainwright, and Michael I. Jordan. DAGGER: A sequential algorithm for FDR control on DAGs. *Biometrika*, 106(1):69–86, 2019.

[39] Eugene Katsevich, Chiara Sabatti, and Marina Bogomolov. Filtering the rejection set while preserving false discovery rate control. *arXiv*, 2018.

[40] J Hu, H Zhao, and Zhou H. False Discovery Rate Control With Groups. *Journal of the American Statistical Association*, 105:1215–1227, 2010.

[41] Ran Dai and Rina Foygel Barber. The Knockoff Filter for FDR Control in Group-sparse and Multitask Regression. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48*, ICML'16, pages 1851–1859. JMLR.org, 2016.

[42] Yongtao Guan and Matthew Stephens. Bayesian variable selection regression for genome-wide association studies and other large scale problems. *The Annals of Applied Statistics*, 5(3):1780–1815, 2011.

[43] Rina Foygel Barber and Aaditya Ramdas. The p-filter: multi-layer false discovery rate control for grouped hypotheses. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 79(4):1247–1268, 2016.

[44] Eugene Katsevich and Chiara Sabatti. Multilayer knockoff filter: Controlled variable selection at multiple resolutions. *The Annals of Applied Statistics*, 13(1):1–33, 2019.

[45] Jelle Goeman, Jesse Hemerik, and Aldo Solari. Only Closed Testing Procedures are Admissible for Controlling False Discovery Proportions. *arXiv*, 2019.

[46] Stephen Bates, Matteo Sesia, Chiara Sabatti, and Emmanuel Candes. Causal Inference in Genetic Trio Studies. *arXiv*, 2020.

[47] J Towns, T Cockerill, M Dahan, I Foster, K Gaither, A Grimshaw, V Hazlewood, S Lathrop, D Lifka, G D Peterson, R Roskies, J Scott, and N Wilkins-Diehr. XSEDE: Accelerating Scientific Discovery. *Computing in Science & Engineering*, 16(05):62–74, sep 2014.

[48] Nicholas A Nystrom, Michael J Levine, Ralph Z Roskies, and J Ray Scott. Bridges: A Uniquely Flexible HPC Resource for New Communities and Data Analytics. In *Proceedings of the 2015 XSEDE Conference: Scientific Advancements Enabled by Enhanced Cyberinfrastructure*, XSEDE '15, pages 30:1—-30:8, New York, NY, USA, 2015. ACM.

A Proofs

Proof of Lemma 2.1. Since $\widehat{\beta}(X, Y; \lambda)$ is a minimizer of the convex objective (6), 0 must belong to its subgradient at this point. This means that there exists an $\widehat{s} \in \mathbb{R}^m$ such that

$$\sum_{i=1}^{n} X^i_{j'} \dot{\ell}(Y^i, X^i \widehat{\beta}(X, Y; \lambda)) + \lambda \cdot \widehat{s}_{j'} = 0 \quad \text{for all } j' \in [m], \tag{18}$$

where $\widehat{s}_{j'} = \text{sign}(\widehat{\beta}_{j'}(X, Y; \lambda))$ if $\widehat{\beta}_{j'}(X, Y; \lambda) \neq 0$ and $\widehat{s}_{j'} \in [-1, 1]$ otherwise. If $\widehat{\beta}_j(X, Y; \lambda) = 0$, then we have

$$X^i_{-j} \widehat{\beta}_{-j}(X, Y; \lambda) = X^i \widehat{\beta}(X, Y; \lambda) \text{ for every } i \in [n], \tag{19}$$

which together with equation (19) implies that

$$\text{for all } j' \neq j, \quad \sum_{i=1}^{n} X^i_{j'} \dot{\ell}(Y^i, X^i_{-j}\widehat{\beta}_{-j}(X, Y; \lambda)) + \lambda \cdot s_{j'} = 0.$$

Therefore, $\widehat{\beta}_{-j}(X, Y; \lambda)$ satisfies the first-order optimality condition in the convex problem (7), so it must be a minimizer. Given the assumed general position of the columns of X and the assumptions on ℓ, the minimizer of the problem (7) is unique [20]. Therefore, $\widehat{\beta}(X_{-j}, Y; \lambda) = \widehat{\beta}_{-j}(X, Y; \lambda)$, as desired. □

Proof of Proposition 2.2. Fix $j \notin \mathcal{A}$. For this variable, $\widehat{\beta}_j(X^{-D_k}, Y^{-D_k}; \lambda(g)) = 0$ for each fold k and for all $g \leq \widetilde{g}_{\text{full}}$. By Lemma 2.1 applied to (X^{-D_k}, Y^{-D_k}), it follows that

$$\widehat{\beta}(X^{-D_k}_{-j}, Y^{-D_k}; \lambda(g)) = \widehat{\beta}_{-j}(X^{-D_k}, Y^{-D_k}; \lambda(g)) \quad \text{for each fold } k \text{ and all } g \leq \widetilde{g}_{\text{full}}.$$

Therefore, the cross-validation errors for the full and LOCO problems match for each $g \leq \widetilde{g}_{\text{full}}$:

$$\mathcal{E}_{-j,g} \equiv \sum_{i=1}^{n} \ell(Y^i, X^i_{-j}\widehat{\beta}(X^{-D_k}_{-j}, Y^{-D_k}; \lambda(g))) = \sum_{i=1}^{n} \ell(Y^i, X^i\widehat{\beta}(X^{-D_k}, Y^{-D_k}; \lambda(g))) \equiv \mathcal{E}_g.$$

Because the rule to choose λ is sequential, we conclude that $\widetilde{g}_{-j} = \widetilde{g}_{\text{full}}$ and also $\widehat{g}_{-j} = \widehat{g}_{\text{full}}$. The conclusion (12) now follows from another application of Lemma 2.1, this time with the full data (X, Y) and the regularization parameter $\lambda(\widehat{g}_{-j}) = \lambda(\widehat{g}_{\text{full}})$. □

Proof of Proposition 2.3. Under squared-error loss, the univariate problem in step 3 of Algorithm 1 reduces to the univariate regression of the residual vector r_j on $X_j - \mu_j$, i.e., step 3 of the procedure introduced in Section 2.1 to motivate the LOCO CRT. Therefore,

$$\begin{aligned}
&p_j(X, Y) \\
&\quad \equiv \mathbb{P}[T_j(\widetilde{X}_j, X_{-j}, Y) \geq T_j(X_j, X_{-j}, Y) \mid X, Y] \\
&\quad = \mathbb{P}\left[-\left\| r_j - (\widetilde{X}_j - \mu_j)\frac{(\widetilde{X}_j - \mu_j)^T r_j}{\|\widetilde{X}_j - \mu_j\|^2} \right\|^2 \geq -\left\| r_j - (X_j - \mu_j)\frac{(X_j - \mu_j)^T r_j}{\|X_j - \mu_j\|^2} \right\|^2 \middle| X, Y \right] \\
&\quad = \mathbb{P}\left[\left(\frac{(\widetilde{X}_j - \mu_j)^T r_j}{\|\widetilde{X}_j - \mu_j\| \cdot \|r_j\|} \right)^2 \geq \left(\frac{(X_j - \mu_j)^T r_j}{\|X_j - \mu_j\| \cdot \|r_j\|} \right)^2 \middle| X, Y \right].
\end{aligned}$$

To evaluate this probability, we need the distribution $\mathcal{L}(\widetilde{X}_j|X,Y) = \mathcal{L}(X_j|X_{-j})$. For each sample $i = 1, \ldots, n$, we have

$$\mathcal{L}(\widetilde{X}_j^i|X,Y) = \mathcal{L}(\boldsymbol{X}_j|\boldsymbol{X}_{-j} = X_{-j}^i) = N(\mu_{ij}, s_j^2),$$

where

$$\mu_{ij} \equiv \Sigma_{j,-j}\Sigma_{-j,-j}^{-1}(X_{-j}^i)^T, \quad s_j^2 \equiv \Sigma_{j,j} - \Sigma_{j,-j}\Sigma_{-j,-j}^{-1}\Sigma_{-j,j}.$$

Furthermore, $\widetilde{X}_j^i$ are independent across i. Therefore, we have

$$\mathcal{L}(\widetilde{X}_j|X,Y) = N(\mu_j, s_j^2 I_{n\times n}), \quad \mu_j \equiv (\mu_{1j}, \ldots, \mu_{nj}). \tag{20}$$

If we let $Z \sim N(0, I_n)$, it follows that for any $c \in [0,1]$,

$$\begin{aligned}
\mathbb{P}\left[\left(\frac{(\widetilde{X}_j - \mu_j)^T r_j}{\|\widetilde{X}_j - \mu_j\| \cdot \|r_j\|}\right)^2 \geq c \,\middle|\, X, Y\right] &= \mathbb{P}\left[\left(\frac{Z^T r_j}{\|Z\| \cdot \|r_j\|}\right)^2 \geq c\right] \\
&= \mathbb{P}\left[\frac{Z_1^2}{Z_1^2 + \cdots + Z_n^2} \geq c\right] \\
&= \mathbb{P}\left[\frac{(Z_2^2 + \cdots + Z_n^2)/(n-1)}{Z_1^2/1} \leq \frac{1}{n-1}\left(\frac{1}{c} - 1\right)\right] \\
&= \mathbb{P}\left[F_{n-1,1} \leq \frac{1}{n-1}\left(\frac{1}{c} - 1\right)\right].
\end{aligned}$$

The second equality in this derivation holds because the rotational invariance of the distribution of Z allows us to set $r_j = (1, 0, \ldots 0)$ without loss of generality. The conclusion follows. □

Proof of Proposition 2.4. Using the resampling distribution (20), we find that

$$\begin{aligned}
p_j(X,Y) &\equiv \mathbb{P}[-T_j(\widetilde{X}_j, X_{-j}, Y) \leq -T_j(X,Y) \mid X, Y] \\
&= \mathbb{P}\left[\left\|r_j - (\widetilde{X}_j - \mu_j)\widehat{\beta}_j\right\|^2 \leq \|Y - \widehat{f}(X)\|^2 \,\middle|\, X, Y\right] \\
&= \mathbb{P}\left[\left\|\frac{\widetilde{X}_j - \mu_j}{s_j} - \frac{r_j}{s_j|\widehat{\beta}_j|}\right\|^2 \leq \frac{\|Y - \widehat{f}(X)\|^2}{s_j^2\widehat{\beta}_j^2} \,\middle|\, X, Y\right] \\
&= \mathbb{P}\left[\left\|N\left(\frac{-r_j}{s_j|\widehat{\beta}_j|}, I\right)\right\|^2 \leq \frac{\|Y - \widehat{f}(X)\|^2}{s_j^2\widehat{\beta}_j^2} \,\middle|\, X, Y\right] \\
&= \mathbb{P}\left[\chi^2_{n,\mathrm{ncp}} \leq \frac{\|Y - \widehat{f}(X)\|^2}{s_j^2\widehat{\beta}_j^2} \,\middle|\, X, Y\right],
\end{aligned}$$

where $\text{ncp} \equiv \dfrac{\|r_j\|^2}{s_j^2 \widehat{\beta}_j^2}$. This concludes the proof. $\square$